PROFESSOR: How you guys doing? I don't know. [INAUDIBLE]. Thanks for inviting us. It's a real pleasure to come here and have a chance to talk to everybody today. I brought one of my senior managers who oversees the [INAUDIBLE] security risks, Dave LaPorte, who's going to talk about some of the more tactical details of what we do. I'm going to talk about the high level.

Feel free to ask questions at any time. So really your opportunity free to ask any anything you're curious about. So there's no decorum in terms of what you guys may ask [INAUDIBLE].

And so I think I was sitting where you guys are sitting-- I don't know-- it's almost 20 years ago now?

PROFESSOR: Yeah.

PROFESSOR: [CHUCKLES] So [? Nicoli ?] and I were a lot younger then. And I was probably a lot thinner by then and had a little bit more hair. You know, one of the nice things about overseeing MIT's infrastructure and operations areas, you can see all sorts of interesting things. And some of the things we'll talk about, a lot of what we do is dealing with interesting problems. And you know, there's no shortage of things in an environment like MIT's.

I think what's really remarkable is we run an open network, which is a little bit of a good thing and a bad thing. We don't have a broad campus firewall, for the most part. And everything's pretty much open.

If you guys want to run a computer in your dorm, right here in the lecture hall, or anything else, you have pretty much unfettered access to the internet, which is,

compared to other schools, actually, fairly unusual. You know, you may not realize that as you sit here, but that's not the norm.

And that brings with it a whole slew of challenges, in terms of keeping things secure. So pretty much, we're wide open to the world. And that means anybody, anywhere, from whatever country, from whatever part of the planet. If they want to reach out and touch your device sitting here in this room, as you're sitting here today, whether it's your phone in your pocket, or your laptop that you're typing on when you're sitting here, they can do that. There's nothing to prevent them from doing that, right?

And that's kind of scary, right? So we did an experiment a couple of years ago. And we just took a device out of box brand new, brand new Apple laptop, and just plugged it in. Registered it for DHCP and just left it there sitting there for 24 hours.

And we left TCP dumps running to just take an inventory of what was coming into the machine for a 24-hour period, just to see what we would see. And then we combined that with, hey, let's go and graph all the various IPs, using GeoIP lookup and graphed them, put them on Google Earth and see what that looks like.

And for one 24-hour period, for an inauspicious or relevant host which is just publicly registered for the internet, it received connections from every country on the earth, except for two. In one 24-hour period, one host, every country, except for two. That's pretty startling, right? Does anybody want to take a guess what the two countries that were not trying to connect to this machine were?

AUDIENCE: [INAUDIBLE].

AUDIENCE: North Korea.

PROFESSOR: Anyone? North Korea is one. Good. Nope. China was very actively represented.

[LAUGHTER]

So might have been the military part, I don't know, but certainly was very actively represented.

AUDIENCE: Antarctica.

PROFESSOR: That's right. Antarctica. Very good. So you get the gold star for today. It's excellent. Yeah. Yeah.

And so, for one 24-hour period, you're seeing yourself subjected to potential attacks, threats, malware, anything else everywhere. For one host. And the entire MIT campus right now compromises about, I'd say, 150,000 different devices. And so, if you do the math-- you can do that outright. You're good at math here at MIT-- that's a lot of threats, right? And you know that happens all day, every day. Right? And that's pretty scary.

And you want to combine that with something else to make a little bit more scared? So Dave and I were sitting in a meaning a couple of months ago. This is kind of a follow-up as to a power outage that happened. Any of you here for the big power outage a year ago or a year and a half ago? It was an exciting time, right?

[CHUCKLES FROM AUDIENCE]

I was here for the big power outage about 20 years ago, when the entire city of Cambridge was out. Now, that was cool. Except it was about 100 degrees, so it was a good time to go over to Boston and see a movie in a movie theater.

But one thing that came out of it that was really interesting to us-- I don't know, let's talk about it for a second-- but [INAUDIBLE] department comes to us and says, you know, this has been really bad, this power outage. We've really had to spend the last four or five months going across the campus and reprogramming all these devices.

OK. Well, they must have some SCADA systems that are connected to their air conditioning, or to the lights in the room, the doors, things like that, right? You would think they do. It's MIT. Makes sense. So sure, you imagine they're secure. Talk about that in a second. And so we figured that was pretty straightforward.

And they said one of the things they had a problem with was their devices keep getting knocked off the net, have these issues. And the more you start talking to them, you peel back layers of the onion. You're like, what do you mean your devices are on the net?

Say, well, yeah, our devices are on the net. Oh. You must have some secure, proprietary control system, a control network. And they kind of looked at us with a blank look and were like, uh, I think that's what it is. That's what the vendor told us.

And this brings up one of the interesting things of the internet of things or the era we're moving into is, for the most part, when I was younger, people using the internet had to be fairly specific, right? You had to fairly know something that you were doing. Today, everybody's using it. And the bar to use it used to be, when you ride the roller coaster, thou shalt be this tall if you're going to ride this rocket. It's gotten a lot lower, all right? And so through the conversation with them, we find out that they have pretty much everything you could think of connected to the internet. Everything.

One of the things that was interesting. MIT launched an energy Initiative, I don't know, five, seven years ago, when Susan Hockfield was here as president. And one of the things [INAUDIBLE] did was really grow the internet of things on campus to create these dynamically managed buildings. So when a classroom isn't in use, lights will go off, heating will change. It's new. They're doing things like that across the campus.

They deployed a gigantic control network. Gigantic. It's actually bigger than our own campus internet is, somewhat. So their control compromises, I think, 400,000 different points that they're monitoring across the campus. Over 75,000 to 100,000 places.

And so the next question you ask yourself-- Dave asks, with the big, bright, wide eyes is, how you guys securing that? They're like, well, we called you guys up, and we were going to check. And they put it in. And they got an IP address by going through a web form and request an IP. Requested IP, it's working.

And you know we're sitting there going, yeah, this whole open internet thing-- how are you guys securing that? And the question, of course, which is, you know, when your blood pressure goes up a little bit was, well, it's secure, because you guys take care of that. So the security's already taken care of.

And you know our look on our face is a little bit-- what do you mean by secure? Well, we have corporate firewalls. It's dealt with. And everything's cool.

[AUDIENCE CHUCKLES]

And you know my next question was, can you show me where that is? I don't know. And you know, of course, the response was, well, everybody does. And of course, we're a little different, going back to the point I made earlier, we operate a fairly open environment.

And we've always believed and it's MIT's philosophy that we believe in defense and security across the stack. You don't want to depend on any one part of the infrastructure to implement security. It's something you have to do at every layer, right? You don't just do it at the infrastructure, you do it at the application. You do it all sorts of places. That's not how these systems-- you know, the internet of things on the SCADA side-- get built. You know, that's kind of scary.

And so one of the things we're seeing here that we deal with, so in addition to dealing with folks like yourself who are doing all sorts of creative and inventive things that keep people like Dave and I up at night. You know, the internet's becoming this utility that's used for all sorts of things across the campus. And it's really changed the dynamics that we have to worry about, in terms of threats, security issues, all sorts of things.

So now, you know, when the internet used to go down or have an issue when we were students, it was an inconvenience. Right? It was annoying. In these most recent types of event, when the internet goes out, people's air conditioning stops working. Your heat goes out, you know? So the threats have really changed.

And so for us, we deal with this whole broad spectrum of things where we operate as a service provider for the campus, providing services to folks like yourself. We also provide services to all sorts of things. And if you combine that with our open network philosophy, it creates a lot of interesting use cases and threats that we have to worry about on a pretty persistent basis.

And people's expectation of how the internet is going to work and one of the things that's also been eye opening is these systems kind of grow. When the power outage happened, one of the first questions they had was, why did the network stop working? Our response was, well, there was no power.

And they said, well, that comes from the batteries on the phone system, right? It's all taken care of. Like, no, it doesn't. That's analog phone technology. Like, well, what's the difference? Well, you guys know a lot.

But there's this expectation that these things operate with the utility service, whether it's at the security level, resiliency level, or everything else. And I love the folks that have so much confidence that these things are being dealt with at that level. But that's a big gap from where we are right now.

And so we spend a majority of our time trying to keep the campus environment running, keep everything as secure as we can. Dave will talk about that in a little bit more detail. He'll give you some interesting stories about the kinds of things we deal with.

But it's an interesting job, right? And we see all sorts of interesting things. And I think the problems just get harder and harder. And I think the thing is is the internet kind of expands.

And the internet of things is a big buzzword. How many of you have heard of that before today? You guys have gone to a Cisco website or something. Trying to sell you expensive equipment.

But this whole phenomenon where everything is internet or IP-enabled, that's here. And unfortunately, a lot of the people writing these systems are not as studious as

people who went to MIT. They create all sorts of interesting problems.

So I think, for us, the real challenge is, when you look at security from a systemic level, it's just there's a thousand little bitty pieces. And it's really, really hard.

And even for us, we have to deal with service providers on the external side. We have to deal with our own customers. We have to deal with application providers. It's this very broad ecosystem of issues you have to worry about to provide security holistically. And the challenge is pretty bumpy at times.

And so from that point, I guess I'll have Dave talk a little bit about a couple of things we've seen. You guys have any questions before Dave gets set up that you'd like to ask on anything in particular? Anything? Come on up, Dave.

PROFESSOR: Oh, you've got a question now.

PROFESSOR: Yeah?

AUDIENCE: Have you seen any APT campaigns attacking MIT directly?

PROFESSOR: Yes. Yes. I think what's interesting is one of the things we're seeing more of is-- you know, one of the things Dave will talk about now that's really hard is visibility. If I told you the story about the one laptop, and we have 100,000 or 150,000 devices, and also, if you think of the number of IP addresses MIT has just having /8 network, finding the needle in the haystack is really hard. So these APTs-- which we do have-- finding that noise or that activity going on within this broad stream of traffic is very difficult.

One of the things we've have help with is some of the tools we have now are a little bit more advanced. And we'll talk about that in a minute. But one of the things we also see now is law enforcement's help from the Federal side or from other parts where they reach out to you to give you guidance in some of these things where they see things like that, which is very helpful.

But you know, operating an open environment, we've had a few of these where I've

been really surprised that they were going on, but they do. And I think one of the other things we're going to see in the future is-- today, we do a lot of research activity here at MIT. It's one of our primary missions. The Federal funding sources that provide for that research don't really have a lot of rules about how you do it.

As you'll find out, as you go into your grad student lives and other things and you get Federal grants for research, whether it's private or primarily from the government, whether it's the NSF, NIH, their requirements are pretty vague.

One of the things we're dealing with lately is they'll say, we have a data requirement in the grant that says, you should have a data policy about all the data you generate from your grant is going to be preserved. All right? And the way MIT does that is it says to the PI, that's great. We have this requirement. Are you going to take care of it? They say, sure. I'll take care of it. Sign the document, right?

Compliance with that? What they did with it? It's left to the discretion of the primary investigator. So if the government was ever to come to us and say, hey, where's the data? Just point to a faculty member and say, hey, talk to home.

But one of the things we're also seeing is the government saying, hey, look. We're investing a lot of money in doing this research. We don't want to spend all the money and give the research to another country-- in some cases, right? So what we've seen on the legislative side or some of the Federal funding agencies side is a lot of them are coming to us and saying, as an industry, I think we some more security departments for this.

And I think what's hard for us is MIT is very much an incubator. We have an incredible number of brilliant people. And for the administration at large, we serve as sort of a hosting company, right? We incubate that activity. We provide them lab spaces, or internet connectivity, or all sorts of things.

But for the most part, it's a fairly federated environment. People have a level of autonomy. And so, as you have more requirements coming up, applying those across the institution is very hard.

To go back to the APT thing, there's a tremendous amount of intellectual capital here. Tremendous amount of interesting things going on here that folks outside this country are very interested in. And I don't know if you guys know this. What country do you think is responsible for more intellectual property theft than most any place in the world? Anyone want to take a guess?

AUDIENCE: That's a dangerous proposition.

PROFESSOR: No, no, no. But I'm serious. I mean, one of the ones that-- it was really shocking to me, because I never expected it. And that's not to say they're doing that here. But anyone want to take a hint? Come on, somebody. Anybody.

AUDIENCE: China?

PROFESSOR: Who? It's not China.

AUDIENCE: Russia?

PROFESSOR: Nope. Wasn't Russia. No.

AUDIENCE: Is it Canada?

PROFESSOR: Nope.

[LAUGHTER]

You're getting close, though, but it's in Europe.

AUDIENCE: [INAUDIBLE]?

PROFESSOR: France. France. That's right.

AUDIENCE: What was it?

PROFESSOR: France. Yeah. So yeah, you wouldn't expect that, right? But I have a bunch of folks who work in the industry side, commercial sector, in the security area. And one of the things they have to worry about is companies, some of them located in this geography, is that's one of their biggest threats, which is kind of surprising. You

would think it's Iran. You would think it's all sorts of other places. No.

And so, it's interesting, right? You wouldn't expect that. And that's not to say the US isn't doing it too. Let's be honest, right? We're just doing it better. We're not getting caught, probably. But needless to say, it's one of the more interesting things.

Another question? You ready, Dave?

PROFESSOR: Sure.

PROFESSOR: Yeah?

AUDIENCE: What sorts of things do you log on an IP sector?

[LAUGHTER]

PROFESSOR: Turn the camera off now. We log some fairly interesting things. I'd say, for the most part, I'll be honest with you, authentication requests, we log. So when you log-in through Kerberos, you log-in through Active Directory, you log-in through Touchstone, through our SAML iVp, those things get logged.

We have very detailed retention policies, which we're happy to share with you. It's published. If you access a web page, if you check in to read your email, things like that, then one of the problems we have is correlating all that information would be fairly hard. That's a lot of different sources. Dave will talk about that in a little bit. But that's pretty much what we do. We try to keep our retention period, usually, within 30 days.

AUDIENCE: So MIT runs its own CA, right? [INAUDIBLE]. Where do you guys keep that [INAUDIBLE]?

[LAUGHTER]

PROFESSOR: OK.

AUDIENCE: Could you use tamper-resistant hardware, or something like that? Or do you just put it on computer somewhere running Linux [INAUDIBLE]?

PROFESSOR: Well, wow. Good question. Yes, we run our own CA. We've been running our own CA since the late 1990s. So in a whole world where the web was doing authentication using username, password over SSL, I think MIT was fairly progressive there, doing PKI.

Now, I don't know. You guys are a little bit younger, but back in 1998, they were telling me the year of PKI is just about here. It's going to be next year. And we've been saying that for about 20 years. So hopefully, soon.

So in terms of how it's stored? Tamper-resistant hardware, that's a nice idea. I like that. That would be great. That's not how we're doing it.

AUDIENCE: [INAUDIBLE].

PROFESSOR: So the CA's usually do it through using USB keys or other tokens that have fairly involved protocols where you can only access information off once it's been written, in terms of you could send something out for signature, but you can't actually get the key off. So they have a whole way of doing that.

To be honest, the CA server we're using now was written before any of those markets existed. And so we've used, I'd say, a typical MIT spirit of being a little bit creative. And so it is stored on the file system. It's stored in such a way to make it difficult to recover. And it's stored in multiple ways. And it's encrypted with multiple keys. And there's a variety of system-specific parameters, which I will not get into every one of them here.

AUDIENCE: So it's more of a security by security.

PROFESSOR: Yes. Yes. So if you know the specific locations, which may not even be files-- I'll give you a hint to look-- and you know how many keys are needed and you know how many bytes to read, you might be able to figure it out. But to be honest with you, if you're able to actually get on the machine, it's game over, right? If you're able to get on the machine, it's game over.

You know, people who tell you that they can still break into the machine and not compromise the key for some ETA, nah. I'd be very skeptical about that. So we try to be fairly secure, but-- anyway, it's not perfect. Other questions?

AUDIENCE: Typically, what percent would you say of MIT [INAUDIBLE] are compromised at a given time?

PROFESSOR: Good question. Not yours, right? I think I'm amazed-- so does everyone know what phishing is?

AUDIENCE: Mm-hm.

PROFESSOR: OK. Good. I can't tell you how many talks I go to where they look at you and they say, isn't that like Seattle? But that's a little bit of an older audience.

Every time one of these phishing things happen-- and they happen quite a bit where these emails go out-- I will tell you, it amazes me. You're the world's smartest institution, all right? As an alum, I'm pretty proud to believe that. You know, we're the world's leading technological institution in the world, right?

I cannot tell you how many people will reply to those things. It just always amazes me how many of them reply to "Dear Help Desk, here's my user name and password. It shocks me, right? And some of them are faculty.

And they call the help desk and go, hey, I wrote back to your quota message. How come my quota hasn't gone up? And oh, by the way, like my inbox is full now. So what happened? Well, they got 200,000 bounced messages to be in their inbox because it's being used to send mass emails.

So I'm being honest. I'd say we see 10, 15, 20 to 30 a month. During one of these phishing spikes, even larger. And I think the ones that have been really interesting are the ones we don't know about, OK?

And the government came to us, I don't know, about a year or two ago and said, hey, we won't get into specifics, but there's a marketplace where you can buy MIT usernames and passwords so you can access library resources. And so they're

bidding on them online on these black markets. If you'd like to access all the materials MIT has in their libraries or on campus, you can simply auction one of these accounts that they've compromised off the web page. And they said, oh, by the way, do you know about this? No. No.

And so we see a tremendous number of those. The success of the social vectors for getting people's credentials is incredibly high, in particular, across our industry-- Dave will talk about it in a little bit, how you try and mitigate those things-- is very high. And it's a little bit scary. Yeah?

AUDIENCE: Dealing on that, is there a way or some web site to see all the places you've logged-in with [INAUDIBLE] or something?

PROFESSOR: Yeah. So we're working on-- one of the things I talked about in answer to the log question was what do you guys collect? We collect a variety of information. One of the things we don't have the ability to do today is to correlate.

So in our case, there's 30 different technology systems involved in some of these things, and different formats, and all sorts of different ways to generate it. We are working to try and make that easier. And our hope is to give the user community something where, from a GeoIP or other standpoint, they can see where their activity is maybe over the last 30 days, or seven days, or whatever their retention period is, to help inform people about where these things are happening.

Dave wants to go a step further. He wants to have you can pick a circle of radius where you're allowed to log-in from geographically. And if you log-in from outside that range, it either shouldn't allow it, or it'll just send a text to your phone to let you know that something happened, which I think is a step in the right direction. But so yes, we're working on that, but we don't have it today.

AUDIENCE: You said there was malicious traffic on the MIT network?

PROFESSOR: Yes. In terms of--

AUDIENCE: What's the primary source? Is it from outside going in? Or is there malicious traffic

from the inside going out?

PROFESSOR: You know, I'd love to say it's completely from the outside coming in. I'd say there's a fair bit of it from the inside going out. I think, realistically, we have a tremendous amount of internet bandwidth and connectivity. We'll talk about some of these more recent UDP reflection attacks, which is a great example.

But when you have big pipes, it's a good resource to use to hurt other folks. And so we see a lot of that, I'd say more so than stuff coming in. Stuff coming in, you see a fair bit, like I talked about with the laptop. But I'd say, in terms of actual volume of traffic, the bigger stuff you see is in and out, in terms of just sheer throughput.

AUDIENCE: How many people connect to [INAUDIBLE] on a given day?

PROFESSOR: On a given day, I'd say like, I don't know, 100,000, 120,000 different kinds of devices. I'd say, people-wise, if you figure people average 2-1/2 devices, probably 35,000 folks, 40,000 folks on a given day. I think what's more surprising is our visitor population's fairly large in a month.

AUDIENCE: What's the policy [INAUDIBLE]?

PROFESSOR: Policy? What? No.

[LAUGHTER]

AUDIENCE: Is there a good reason not to?

PROFESSOR: So you know, MIT's a very open place, all right? And I think that's one of the great beauties of being a student here and one of the things I've always cherished about being here is we're a place where it's OK to experiment. It's OK to do things. It's OK to learn about things. It's OK to develop new things.

That's one of the great things about being at MIT and it's what's special about being here too, right? That's what's pretty unique. You don't need to go to the policy office to say, hey, I want to run a core exit node today, or I want to invent a new anonymous protocol, or something like that. That's one of the thing's that's really

unique about working here and going to school here.

And I think, for us, is it a good idea or is it a bad idea, depends how you're trying to do it, right? If you're doing it as part of some thesis research into an automatization technique, some privacy, it's probably fine. If you're doing it for the purpose of running some kind of black market ring or something like that-- I mean it's probably not a good idea.

But from a policy standpoint, MIT's fairly flexible. We really try to balance the need for-- the institution has a responsibility to behave responsibly, right? Let's just be honest. As an institution, we have to do that. But we try as much as possible not to encumber the activity of innovation. And so, for the most part, that's worked out pretty well. I'd say MIT's been fairly successful over the last 125 years.

But you know, I think it's one of those areas where if one of those activities was to place the institution collectively in jeopardy, then we have to look at that. But MIT does run a variety of [INAUDIBLE]. SIPI has some. CSAIL has a few. They show up on Dave's naughty list like a plague, but we do that. You can't run that at most schools. Questions?

My esteemed colleague, Dave LaPorte. He from used to work in the Harvard University network. And he'll talk a little bit about networking at a liberal arts school, if you get him outside of the office. And he's also a teacher himself at Northeastern [INAUDIBLE]. So Dave will talk to you a little bit about some specific examples of the kings of things that keep us up at night.

PROFESSOR: Oh, wait. Wait. I think I got it. Loose connection. OK. [INAUDIBLE].

All right. Hello, everybody. My name is Dave LaPorte. As this very verbose title claims, I am Manager of Infrastructure and Security Operations, which basically in a nutshell means I'm responsible for maintaining and operating and securing mit.net, which is definitely a full-time job.

Today, I've got a lot of content to cover. I'm going to cruise through it, to leave room for questions at the end, because I think that's probably where you guys get the

most out of it. If anything's not clear, just stop me. Raise your hand. I don't mind being interrupted.

But yeah, with that, we'll start talking with the Security Operations Team, which is really the one central body that does security as a full-time job here at MIT. We'll talk about some of the events that we've had in the recent past here and what we've done to mitigate them and to leave the current state of security at mit.net.

We'll talk about the current landscape, what we're facing a lot of now, which-- as Mark alluded to-- is to a large degree, social. And we'll talk about some future trends, which are nebulous by nature, so the slides are kind of sparse.

So the team there, as you can see is Mark. I report directly to Mark. Under me, there's a Security Operations Team led by our team lead, Harry Hoffman. He has three people under him, Andrew [? Munchbach ?], who basically is the analyst who does a lot of the washing of the systems, does a lot of the notifications to users, responds to complaints from the outside world. Then Mike Hossle, who does a lot of the engineering activities, a lot of the forensics. And you have Monique Buchanan, who handles a lot of the correspondence and community outreach. Harry himself is also extremely hands-on.

So I just want to preface this whole thing with we have a team of four in a very large institution with tons of devices. So the federation that Mark talked about is really a necessity, in order to even try to secure a network of this size.

So now, the portfolio services we have and what-- we're going to blast through this stuff fast-- consulting, we talk with people and help them on campus. Services, we provide some services to the community. And the tool set that we use.

The services we provide are pretty varied. We do abuse reporting. So this is response to complaints from the outside world, typically, the vast majority of which are Tor exit node-related.

[LAUGHTER]

They just are.

PROFESSOR: [INAUDIBLE].

PROFESSOR: Endpoint protection, so there are some tools and products out there that we install on both the community at large machines-- you opt-in if you prefer. If you're part of the MIT domain, which is typically administrative staff, some might be auto-installed for you.

Network protection, these are tool sets that we have, either at the border or throughout mit.net that detect anomalies or capture flow data for analysis. Data analytics helps us correlate, put all this stuff together and try to get some actionable intelligence out of it. Forensics are-- well, we'll talk about those in a second.

Risk identification, basically probing and assessment tools, basically Nessus and things that look for PII, Personally Identifiable Information. which, being in Massachusetts, we need to comply with 201 CMR 17.00, which is a Mass regulation that requires us to be able to identify where all the PII on our network lives. Outreach awareness and training, just what it says.

Compliance needs, this is, in large part, PCI DSS. So PCI, being the Payment Card Industry, has DSS, which is the Data Security Standard. Believe it or not, MIT-- well, you'll probably believe it-- MIT is a credit card merchant. We have multiple vendors on campus, and we need to be able to make sure that that infrastructure is compliant with PCI DSS. So Security is a part of the team that basically manages and ensures that compliance.

PCI 3.0, which is the sixth major update to the standard, goes live on January 1. So we're kind of in the process right now of ensuring compliance of all of our infrastructure. And providing reporting alerting metrics on the work we do.

So here's some of the end point protection products we use. This eagle-- I think it's an eagle there-- is a tool called CrowdStrike, which is currently being tested within IS&T. Basically, it's a tool that watches for anomalous behavior from a system call perspective.

If you're using Word, and Word suddenly starts doing something that it shouldn't do, like maybe trying to read the account database off of the system and a bunch of passwords, it alerts and throws a flag. It's a cloud-based tool, which we'll talk more about later. So all this data gets sent to a central console. And should machines start doing things untoward from a heuristic or behavioral perspective, they get red flagged.

GPO here, these are just Group Policy Objects, so managed systems which push down policy. The S is Sophos. It's anti-xrays, anti-malware, anti-spam-- oh, not anti-spam-- but anti-malware, anti-virus, all of the typical stuff that we expect when we buy and end point protection product. PGP does hard drive encryption for select systems on campus that have sensitive data.

Some of these tools are in flux. The industry seems to be going more towards a more vendor-neutral solution, if you want to call it that, so BitLocker on Windows, FileVault on the Mac. So we're exploring those options as well. And Casper is a way to manage, mostly Macs, to enforce policy on managed Macs.

On the network protection side-- I'll just start down here. Akamai is a company that came out of MIT, has a lot of MIT alums. They also have extremely good services, so we have partnered with them on a lot of their services. And we'll talk about them fairly extensively.

TippingPoint is an IDS vendor, an intrusion detection system. As I said, some of these tools are in flux. That might be one of them. But we basically have an intrusion prevention system at our border. We don't actually prevent, we just detect. So we don't actually block anything on the MIT border, except for some very basic anti-spoofing and standard rules you'd find anywhere.

Stealth Watch is a tool that generates NetFlow data-- or I should say collects NetFlow data. So we use Cisco devices, but all network devices will output details, meta information about the flows that they're sending, source port, dest port, source IP, dest IP, protocol, et cetera. StealthWatch collects this, does some basic security

analysis on it, and also provides APIs that we can interface our tools with to do some more intelligent things.

And RSA Security Analytics, this is another tool. [BANG] Oops. It's, in a lot of ways, like an IDS on steroids. It does full packet capture, so you can actually see some content if things get red flagged.

On the risk identification front, Nessus, kind of the de facto vulnerability assessment tool. So we typically use this on-demand. We don't unleash this on 18/8 at large. But if we get an on campus DLC that would like us to come in and perform some basic assessment for them, we can use Nessus.

Shodan is a computer-- they call it a computer search engine. Basically, they scan the internet at large and have lots of, lots of good security data. We have a subscription, so that we can leverage that intelligence. And Identity Finder, that's a tool that we use in locations where there's PII, Personally Identifiable Information, in order to comply with mass regs and just to make sure we know where critical data lives.

Forensics is a business that is-- periodic? I'm looking for the right word. This isn't something we do, until we do a lot of it for a long time. And then we don't do a lot of it. Basically, when cases surface, we have the tool sets. EnCase is a tool that allows us to image drives and go through them looking for content. FTK, the Forensic Toolkit and the Sleuth Kit are other tools.

We often get called in for cases where we have to image drives for intellectual property cases or whatever cases the OGC, the Office of the General Counsel, needs to have computers imaged for. So we have all the tools that's necessary to do that. But frankly, it's not our day job. It's something that comes up occasionally.

So how do we put all this data together? Mark alluded to correlation. We have operating system logs for managed systems. You can see that we have NetFlow. We have some DHCP logs, IDS logs, Touchstone logs.

Splunk is a tool that does a lot of this correlation work and take data that's not

necessarily normalized and normalize it and allow us to correlate across different sources to get more intelligence. So when you were talking about-- whomever out there asked about, maybe, a log-in page that could show where you last logged in, et cetera, Splunk would probably be the enabling technology for that, because we can put everything together, we can do GeoIP lookups, and really build something on top of the raw data, pull some actual wisdom of that data, so that we could present it to you in a page.

OK. So now, we'll talk about attacks and things that you might find more interesting. We'll talk first about Distributed Denial of Service attacks, which we've really received a lot of in the past few years. We'll also talk specifically about attacks that resulted from the Aaron Swartz tragedy of a few years ago, which ties in to the Distributed Denial of Service attacks.

OK. So just a primer on Distributed Denial of Service. I apologize if this is remedial. So Denial of Service attack really attacks the A of the CIA triad. CIA triad's the foundation of computer security. It's Confidentiality, Integrity, and Availability. So we're going after the availability, right? We want to take a resource down so that legitimate users can't use it.

That could be defacement of a page. Very simple, right? Digital graffiti, just ruin the page so nobody could see it. Could be resource consumption where you eat up all the computation on a system, all the bandwidth on a network. Could be a single attacker. But much more likely nowadays, you're going to invite your friends and you're going to have a party and a DDoS, a Distributed Denial of Service.

OK these are recent trends in the industry. These are pulled from the Arbor Networks State of the Internet report. Hacktivism is the most common motivation. According to them, it's 40% of all claimed-- actually, those attacks that are attributed, 40% of them are attributed to hacktivism. The next one is 39% unknown. So it's dominating the top of the heap.

Last year-- and these numbers, I believe, are from 2013-- there were multiple 100Gbps attacks. So the year before that made news because I think there was an

attack on Spamhaus that was 300Gbpm. This following year, that's kind of-- I wouldn't say the norm, but we're seeing a lot more of that.

Longer-lasting attack, this operation here, Ababil, was a multi-month attack against the US financial sector. It went on for months. It was 65Gps sustained at times. I've heard stories about this-- but you can Google more online-- where it was just relentless. They just couldn't stop it. And we'll talk about the way that they ended up doing it.

But frankly, at 65Gbps or at 100Gbps, you're at the mercy of the attacker, right? There's very few organizations on the planet that can sustain an attack-- can sustain-- can survive a sustained attack of that magnitude. You just can't do it.

And we're seeing a shift towards reflection and amplification attacks. So this is where you take a small input and generate a large output. This is nothing new. This goes way back to-- let's go ahead a sec, all right, I left that slide out-- but this goes way back to the ICMP Smurf attack where you would ping a broadcast address of a network, and every machine on that network would respond to the supposed originator of the packet, which, of course, would be spoofed, right?

So I would masquerade as Mark. I would sent a packet to this class's broadcast address. And you would all respond with packets to Mark, thinking that he sent it. Meanwhile, I sit in the corner and laugh. So this is nothing new. This goes back-- I mean, when I was in high school, I was reading about this stuff.

[LAUGHTER]

So UDP and ICMP, but UDP is a fire and forget protocol, right? It's not TCP. It's not reliable. It's not connection-oriented. So fire and forget, it's easily spoofable. And over the past year, what we've seen are exploits of amplifiable features of these three protocols, in particular.

So DNS-- this isn't working. It works with a clicker. DNS, port 53, UDP, right? Basically, if you sent a 64-byte ANY query to a misconfigured server, it would respond with a 512-byte response. So that's an 8X amplification factor there. Not

bad.

What we found personally here on mit.net was, when this whole trend started, like most things-- EDUs and particularly here-- we were at the forefront into this trend. We were seeing this before it really took off against commercial victims. But we saw a 12-gig DNS amplification attack here, which substantially impacted our outbound bandwidth. We have sufficient bandwidth. But at those rates, if you add that to legitimate traffic, we started to notice issues. And Mark and I had to come in and resolve that.

SNMP, which is UDP port 161, very useful management protocol. But if you send a GetBulkRequest of 64 bytes to a device that's improperly configured, it will respond with up to 1,000X amplification. So that's even better, right? If you're an attacker, you're going to target things. And we saw huge attacks against printers on campus. So they would have a printer with an open SNMP agent. They would send packets to it. And we would send back 1,000 of them and pollute the internet.

NTP, Network Time Protocol. In this case, a misconfigured server would respond to MONLIST command, which would-- I'm not sure of the amplification factor on that one. But that one's really, really popular. So we got hit pretty hard with the NTP MONLIST misconfiguration.

And we ended up doing a few things to mitigate all these attacks. So on the NTP side, we disabled MONLIST on the NTP server we could, which kind of kills the attack in its tracks. But this being a federated institution where we don't have power over nearly anything-- I shouldn't say that-- nearly everything, there's just a lot of things we just don't have the reach to touch or the authority to touch.

So what we ended up doing was just rate-limiting NTP at the border. And that's been in place now for almost a year? Almost a year with almost no negative impact. So we rate-limited down to, let's say, a few megabits, which was certainly better than the gigs we were sending out to the internet previously. So that's kind of a solved problem.

DNS is a bit harder, was a bit harder to take care of. What we ended up doing was started to leverage an Akamai service called, eDNS. So Akamai has this service where you could host your zones with them. They're one of many providers. But we had an existing relationship with Akamai, which I'll talk about in a minute.

So we leveraged their eDNA, bifurcated our DNS, our domain name system space. We put an external view on Akamai. We put an internal on the servers that always served MIT. And then we ACL'd off the internal view, so only MIT clients could hit our internal servers. And the rest of the world hits Akamai.

The benefit of Akamai eDNS is it's hosted in a content distribution network. It's all over the world. It's being served out of Asia, Europe, North America East, West. It's all over the place. Most people can't take down Akamai, so we don't have to worry about our DNS going down any more. So that's kind of how we resolved that problem.

OK. So these are details of the attacks themselves. Source obfuscation, this is probably remedial. Why you do it, to avoid detection and prosecution. I'll skip that one.

OK. So maybe you don't want to hide your address or spoof your address, you just want to destroy a target with bots. So botnets are huge right now. The "it's OK, no problem, bro" was used in that operation Ababil, which really targeted the US financial sector.

So in this case, rather than just spoofing a bunch of packets from one host, we're using a botnet of legitimate systems that don't necessarily need to just spoof. Since these are legitimate systems and they'll respond to, say, a TCP synack, we can actually do more higher level attacks, like attack an HTTP server and do GET and POST floods.

They might hire stressors. Stressors are basically botnets for hire where you hire them to do load testing, and they go and load test someone else for you. There's, no doubt, probably legitimate ones out there. But there are others that aren't and

are basically denial of services for hire.

OK, so the mitigation strategies. We talked about one, which is DNS. We can use DNS to mitigate these attacks. So we have used Akamai to do that. In this graphic here-- which is probably too small for you to see-- but basically, this slide is way too far ahead where it should be.

So OK, we had an attack against our web server, so I'll just brief you real quick. One of the attacks that followed the Swartz tragedy was an attack against web.mit. They took down our web server.

The way we solved that was we used our bifurcated DNS to point internal clients to web.mit internally. And then we use the Akamai content distribution network to basically mirror web.mit. And then we used the external view of our DNS to point external clients to Akamai.

So when a user out on the internet-- which we'll say is over here-- wants to go to web.mit, they actually go to the Akamai CDN, which serves up the content. If it's content that, for some reason, they can't directly serve out cache, it's dynamic or whatever, the origin server, which is still web.mit, and Akamai will go and fetch the necessary content, send it to the user, and then potentially cache it for some interval. So short story here is that the attack-- I'll talk about it in a minute-- but the way we solved it was we put the actual web server on the content distribution network of Akamai.

The other attacks that we'll talk about-- this how we mitigated. These two slides are out of order. So I mentioned a few attacks. I mentioned the NTP attack. I'm going to mention a couple of others. But basically, these are attacks that are just brute force, trying to overwhelm our bandwidth.

And I mentioned, when you get up into the tens of gigabits range, a lot of internet end users, such as MIT-- you can name another service provider, but a very large user-- even we would have trouble handling tens of gigabits of traffic. So in that case, your options are really limited, right? If it's spoofed traffic, how do you put a

filter at your border to block this traffic? And even so, once it's got to your border where you filter it, it's already flooded your pipes. So how do you do this? You have to push it back up into the cloud, into the internet, and block it there.

And the way that many people are choosing to do it and the way we've done it here is through BGP mitigation. So if you're familiar with BGP, which is Border Gateway Protocol, it's the protocol that runs routing on the internet. And it's a path vector protocol that uses ASN, so autonomous system numbers. So every multi-homed organization on the internet has an ASN, an autonomous system number. And BGP uses that number to build paths through the internet so that you can have multiple paths to get to a particular ASN.

In this case, I'm using example 123, because I created this for another organization. We're three, because we're awesome like that. Harvard was 11, so they're a little slower to the punch.

But in this case, we've got a path. So we've got the beginning of the path is ASN123. The end of the path is 789. And there's some sequence of ASes, autonomous systems, that this packet has to pass through. So what we're going to with BGP mitigation is just inject another ASN into the mix. And that ASN has the capabilities to handle this traffic on our.

So in this case, we have ASN456. And they are going to be kind of a sanctioned man in the middle for us. We're going to allow them to advertise our prefixes so that, when we come under attack, if 18.1.2.3.0/24, a small slice of 255 addresses at MIT comes under attack, we allow this AS456 to advertise that prefix on our behalf.

Once that change propagates across the internet, all of the traffic starts going into that AS. And in this case, for us, that AS is Akamai. And they have lots of scrubbers and can handle the high bandwidth that we can't.

So on the back end of that connection is a private connection we have into Akamai where they send the post-scrubbed, the clean traffic out to us. And that way, we can avoid these sorts of potentially deadly attacks that could just take us offline. If you're

getting hit with that much traffic, there's just nothing you can do.

So actually, before we keep going, any questions on what we've covered so far? Yes?

AUDIENCE: This is just more of a networking question. You mentioned the borders.

PROFESSOR: Yes.

AUDIENCE: And [INAUDIBLE]. So I'm trying to understand the structure of the micro-network and just what a border actually is.

PROFESSOR: Let's see. Let me see if I can pull up a quick diagram for you.

PROFESSOR: [INAUDIBLE].

PROFESSOR: Oh, that's right. Yeah. It's a video camera, it's not--

So MIT really has three border routers, external 1, external 2, and another router. Let's just call it external 3. So these are our-- so basically, the actual mit.net is pretty much a standard hub and spoke topology. We have core switches connected out to a distribution layer. And then they go out to access layers, which are basically buildings.

At the border, we have these three borders here. This is incredibly vague. I apologize. But we have multiple providers. So for instance, our commercial providers will soon be dual home to both of our border routers. This external router 3 here-- which is not its real name-- but this external router here is basically for research peering.

So we kind of have a delineation between commodity and commercial peering and research peering. So all of the BGP we're talking about happens between the external border and our providers and back up into the internet itself. Oh, and these are just choke point routers that we have between our border and our core.

OK. So in response to the Swartz tragedy of-- I believe it was two years ago, I had

just started here-- certain hactivists took it upon themselves to attack MIT as an institution. So we experienced three attacks. And I'm going to go through all three, because there were three separate and distinct types of attacks.

So the first attack was against our infrastructure itself. So at the time, MIT did and does and will support openness. And we have a very open network, especially in comparison to other dot edus, having come from another one. That can be a blessing and a curse from a security perspective, right? So we're open to the world.

In this case, our border routers, these guys here that we just drew, were running an older version of software that was vulnerable to a particular denial of service attack. So the attackers in this case sent a very low bandwidth stream. I mean, it was really low. It was less than 100k. It would have been totally non-noticeable, without actually going on the device and debugging.

They send it to the management interface of those devices. And those devices promptly just keeled over, right? They didn't die, but the CPU spiked. They weren't routing packets. mit.net, at that point, was offline.

So in this case, this was the first attack we experienced. I think it was during the Patriots playoff game. Sometimes, I think these things are planned to find it when staff is not paying attention. So it was during that playoff game.

What we ended up doing was immediately upgrading the software to a patched version, right? That was a quick triage fix. The longer term fix was that-- outsiders on the internet probably don't need to access our management interfaces, right? A very select few need to access those interfaces. So we ended up implementing, basically, the least privileged, so that only at the IP addresses of our staff on vCAN could access them.

And we stopped using clear text management protocols. So that one was fixed. Then attack two came in.

PROFESSOR: There's a question over here.

PROFESSOR: Oh, a question. Sorry.

AUDIENCE: So [INAUDIBLE] that this was [INAUDIBLE] attack against the service provider?

PROFESSOR: I think it would be fair to say it was not a [INAUDIBLE] attack.

The second attack was against web.mit.edu itself. And this was what I alluded to on the DNS mitigation slide that I got ahead of myself on. So web.mit was in our data center, protected by a firewall. So it was behind a firewall.

What ended up happening was that the attacker sent a flood of HTTP traffic. It was a GET and POST flood. I'm not sure which one it was. But basically, they didn't kill the web server, they killed the firewall. The firewall keeled over because firewalls too are a blessing and a curse, right?

A stateless router access list is very simple, is very fast, but you also lose a lot of the granularity in what you can filter. Because it's doing it packet by packet, you can only use the criteria in each packet, ports, and IP addresses, mostly. We hid behind a firewall, which worked well when it worked. But when it came under load, the state required of the firewall-- because of firewall tracking in every state, in addition to the packets-- it died.

So the triage fix for this attack was that we moved it to a routed network. And that's something we would have preferred not to do, but you really had to, due to the attack that was ongoing. The longer term mitigation that we performed was that we moved into the Akamai CDN. So you may notice, if you go outside of MIT, as you go to web.mit now, it doesn't resolve to 18.09.22 anymore. It resolves to a C name, which in turn resolves to an Akamai IP address.

And attack number three, this one actually wasn't on the side of mit.net. This was on the side of our registrar. So we found the homepage of MIT-- www.mit and web.mit replaced with this page here. And we quickly did some diagnostics on the web server. Everything looked fine. The server was not compromised. It was not defaced like this.

And what we did end up finding was that our who is information for our name and our actual DNS delegations weren't working. So in this case, you can see the administrative contact, "I got owned." And then our address, "Owned network operations, Destroyed, Massachusetts." They were clearly just trying to poke at us. But it was delegated out to these two servers at CloudFlare, which is a cloud-based hosting provider.

So this is what I call the troll. This was on Gizmodo. This was a bit of indirection on the part of the attackers, however many there were. The hack went down like this. So this is what he told the world very soon after this happened.

So once we realize that it wasn't MIT or anything on mit.net that was hacked, it was actually our registrar, we got in contact with our registrar. We got our records changed. We got everything locked down.

But of course, in DNS, there's time to live values involved. Some of them are hours. So after this attack was resolved, there were still some flux afterwards. And during that time, we're trying to clean everything up, he posts on Gizmodo in the comments on an article about this. "Own the MIT NOC guy with a browser exploit." That guy. "Get their Educause logins, which were blah." So he goes through the whole scenario. But the vector here was the MIT NOC guy.

So Mark is swearing at the time, up and down-- not actually swearing-- but swearing up and down that it's not him. He's not compromised.

So what we did end up finding out after the fact was this was published well after the incident. That link is still live, if you want to read about it. This was indirection. It wasn't true. The actual vector was that our entire registrar was owned. Our registrar, being .edu, is run by an organization called Educause.

Turns out that every DNS registration account had been compromised. The attacker had this for I don't know how long, but they just decided to show their hand with the MIT hack here and actually use their power to expose themselves, but to also hack our DNS. So this one wasn't actually anybody at MIT's fault, it was on the part of our

registrar, which they soon acknowledged. This was in February of 2013. They mentioned that they were in fact breached. And we all had to change our passwords. And do we have two-factor now?

PROFESSOR: No.

PROFESSOR: OK. But we ended up locking down our domain account so that it couldn't be changed. But you know, it turns out, if the entire system is owned, if you check a box that says, "locked," which prevents people from updating it, it doesn't do much good. In any case, they've fixed their system. And this one wasn't our fault, but it was kind of an interesting one because it really subverted some of the core protocols of the internet in order to do this.

OK, so the current threat landscape. No, no, no. [INAUDIBLE]. Ah. So if you can't exploit the silicon, exploit the carbon. Exploit the use at the keyboard. And this is what we're seeing a lot of now.

I mean, from my personal experience, having been in this for almost 20 years, the network-based vector-- with the exception of this year, which I'll talk about in a minute-- but the network-based vector, attacks that originate on the network and remotely exploit hosts, that isn't a lot of what we see anymore. Computer systems actually do seem to be getting more secure, from the outside at least.

You know, Windows, and Solaris, and Linux, in the old days, maybe a decade ago, they used to ship with all their services enable. I called it lit up like a Christmas tree. Everything would be on, because on the convenience and security continuum-- if we agree that one exists, some people don't-- but they were way on the convenient side. They wanted everything to work out of the box.

Whereas, I think we found a sensible medium where, when you install a fresh operating system, there is a host-based firewall running. And there aren't world accessible services running on a system. We've also got things like Windows Update and Apple Update and package managers in all the Linux [INAUDIBLE], so that a box that gets online, pretty quickly, will get itself up to date. So you don't have

these ancient boxes with ancient services open to the world.

OK. So where I was going with that is they've moved up the stack now, right? Maybe they're at level eight or nine now. They're dealing with people. And they're trying to exploit human failings or frailties, like fear, or greed, or trust, or familiarity, to kind of leverage credentials, so that they can exploit application access or privileged access. Rather than exploiting hosts themselves, they're exploiting people.

A few things we've seen on campus in the very recent past. This one here-- oh, I should say, we have not experienced this one on campus. I don't want to scare anybody. This has been seen, though, across the nation in different .edus. And it's a very serious threat, so we're moving quickly to address it. However, we have not actually seen it here.

So this threat is spear phishing. And this is probably remedial, but spear phishing targets a particular community with a plausible message. So if you just get spam messages or phishing messages, they're just casting a wide net and catching people who, for some bizarre reason, might respond to that.

However, in a spear phishing attack, if they're able to narrow it and find a community of interest, that they can actually say something that sounds mildly plausible, Bank of America customers, or MIT students, staff, and faculty. In this case, they were able to choose communities of different institutions. I'm not singling anybody out, but it is public. And Boston University was one of them.

They targeted the community with bogus email messages pointing to a bogus authentication site. Some percentage of the users actually clicked through and logged in, which of course, gave the attacker the credentials. The attacker then went to the legitimate site and redirected those user's direct deposits to an account under their control and emptied that account. I'm not sure the dollars affected there are public, but it was probably a large amount.

So in response to this-- I mean, how do you combat that? To a large extent, it's a

user education issue, so community awareness. Just let people know you can't trust email. You just have to verify things, before you click on them.

But from experience and just knowing human nature, that's only going to get you so far. There's always going to be some percentage of people who click on it. And I'm amazed. As you can imagine-- you all probably are as well-- I'm a repository of all the phishing messages my family receives. So I'm getting stuff from my father, my sister, like, is this legit? And it's getting increasingly hard to tell. So you've got to actually go into the headers and see where it's coming from.

And a lot of male clients now, they don't actually like to show you the link that it points to, which gets really annoying. So it's getting harder. So some of it is just user error, for lack of a better word. And some of it is it's just increasingly difficult to tell.

So at the root of the problem, in my opinion, is that passwords are just a dead technology, right? In terms of the factors, it's something you know a password, it's something you are, maybe a biometric, or something you have, a token. So what we're doing to try to mitigate this here-- because basically, these man-in-the-middle attacks are just stealing something you know. They're stealing your password. Well, if we can also add something that you have, that attack is only going to get them halfway. They're not going to be able to compromise your identity.

AUDIENCE: [INAUDIBLE] number?

PROFESSOR: You won't be able to compromise-- the attacker won't be able to compromise the identity of the user. So we are rolling out a second authentication factor in the near term that will be tied into our Touchstone IDP. We can release that to these guys, right?

So early release, if you're interested. If you're go to duo.mit.edu, we're using a vendor called Duo Security. It's a cloud-based two-factor authentication system that's being used in a lot of edus. But basically, you'll register your phone as a second factor and you run a little smartphone app. If you don't have a smartphone, I'm sorry. But if you don't have one, you can also do it via SMS. It will actually call

you and tell you a number. And you can also generate one-time passwords, a list of 10 passwords that you can use.

So this is coming soon to the community. It is completely active at this point, it's just not announced. If you want to go to duo.mit, you can opt your phone in. You can actually turn on Touchstone. I'll give you a quick demo, to show you how it integrates.

The beauty of using standards and federated systems, like SAML and Shibboleth, which underlay Touchstone, is that we can easily lay on additional factors. So in this case, I'm going to just go to a tool that I use, which is transparently going to authenticate me using Touchstone with my certificate. But I should now get prompted for-- too fast there. Who runs this now?

PROFESSOR: Don't know that guy.

PROFESSOR: I don't know.

[LAUGHTER]

PROFESSOR: He got owned.

[LAUGHTER]

AUDIENCE: Yeah.

PROFESSOR: Give me a site to hit. I'm just drawing a blank, because I'm in front of a crowd.

PROFESSOR: Splunk?

PROFESSOR: Splunk. No, not Splunk. Because that's not a native Touchstone integration. Atlas.

PROFESSOR: Atlas, yeah.

[LAUGHTER]

PROFESSOR: All right. All right.

PROFESSOR: Well, he's very secure, at least, right?

[LAUGHTER]

The only real security is to stay off the internet.

PROFESSOR: All right. If this doesn't work in 20 seconds, I will move on.

PROFESSOR: Do you have another browser, maybe, that's cached?

PROFESSOR: Maybe. There we go. Oh, weird.

OK, so here's my typical Touchstone login. Now, you're going to just get one more prompt.

PROFESSOR: You're actually off the Wi-Fi, it looks like.

PROFESSOR: Oh, yeah. I did. All right. Well, let's leave that demo for after class, and let's edit that out.

[LAUGHTER]

PROFESSOR: In fact, now I tried it--

PROFESSOR: All right. I did not rehearse that demo, so that's what I get. But trust me, it works. If you go to Duo, you can register yourself. And all of your two-factor interactions-- or I'm sorry, your Touchstone interactions will all be two-factored now. And you'll be super secure. It really does work well.

Another threat we've experienced in the past few months is-- this, again, is something that's targeting not just edus, but organizations across the country. But we've been getting police called ID spoofing. And this kind of transcends the digital world, for the most part. Members of the MIT community are getting calls from local police departments nearby their hometown.

PROFESSOR: Yeah. It appears to be a call.

PROFESSOR: Yes. So these police departments are telling them bad news, right? They're telling

them you're about to be charged with a crime. I think some of them are tax fraud. Your family member's been in an accident. Of course, it's not real. Their call's coming from an attacker who's using ANI, which is automated--

PROFESSOR: They're using SIP, basically, to forge their From field in their voice call, right?

PROFESSOR: OK. Which feeds into caller ID.

PROFESSOR: Which then gets translated by a bridge who trust the From field that a SIP message to be whatever number's placed in it.

PROFESSOR: OK.

PROFESSOR: And so the attacker's like, all right, so I sign up for a cheapie SIP service. And I'll set my From field to be the police department's number in Lexington. And I'll send it.

And once it gets to the transcoding gateway that turns it back into traditional telephony, it says, all right. Well, that's the number it's from. We'll just show it to the user when it shows up.

PROFESSOR: So you end up getting a call from what you think is a police department with extremely bad news. Again, they're exploiting human frailty here. In this case, it's probably fear. Maybe it's a little anger, but it's the case--

PROFESSOR: Or guilt.

PROFESSOR: What's that?

PROFESSOR: Or guilt.

PROFESSOR: Yeah. Yeah, I suppose. [LAUGHS] If it's legit.

But in any case, these calls all come to the same point, which is you need to pay a fee for something, which, on its face is kind of bogus. But once they've told you this-- I mean, if they told me that my wife was in a car accident, I would not be in the right state of mind. And I might believe some craziness that they tell me after the fact.

So we've had people on campus who-- I don't think it's actually anyone's actually paid, but we have had multiple targets of this attack. And again, this isn't just here. If you Google it, you'll find reports of it in Pittsburgh newspapers and all over the country.

But again, this is a spear phishing attack, right? They figure out where you live. And they can do this by just Googling your name or just looking at MIT and checking the directory and finding out some details about you. And they just need enough to be plausible. And then they proceed to call you and scare you and try to extort you.

The mitigation here is difficult, because it involves the phone system. It involves multiple bridges. Frankly, I'm not a phone guy, so it involves lots of stuff I'm just going to wave my hands about. But law enforcement needs to be involved here. You have to go to your providers and work with law enforcement to get that traced back.

Actually moved forward? And one more along the spear phishing email side of things. There's another thing that some people call whaling, which is the intentional targeting of high level staff at an organization.

We have experienced this-- I've experienced it elsewhere, but we've also experienced this here where they will send extremely targeted messages at high level staff, using org charts, using directories to find plausible reporting relationships, and write some really believable messages that, if you're not careful when you click reply, that wire transfer number that they're asking for isn't going to the person that you think sent the message. It's going to somewhere on the other side of the world or the other side of the country.

So we've been experiencing this as well. The short of it is-- in your own security class, you know this-- SMTP is not a reliable protocol. You can't believe pretty much anything in there, which just runs completely counter to human nature. So take everything with a grain of salt.

PROFESSOR: You going to tell the story?

PROFESSOR: No, you can.

PROFESSOR: Yeah. Yeah. So you could tell it. Basically, we had one recently. And I won't go into the details or the names. Keep anonymity.

But a senior member of the administration reaches out and says, hey, I got this email. This is from someone and says, I need help-- a very senior executive-- with a wire transfer. And so he sent me this email. And I replied. And he said he didn't know what I was talking about. And so how did that happen? Is my email account hacked? It says I sent this email, but I didn't.

And to Dave's point, the whole possibility that email itself could somehow be spoofed, without your account itself being compromised, is very foreign to people, right? It's a trust relationship.

So it turns out, I guess, it was someone at an internet cafe in Nigeria or something like that, which we were joking about. But yeah, they basically went to the MIT org chart, found a senior executive, found someone in the Vice President of Finance office and said, hey, I need you to help me with this wire transfer call. Here's the number. And these are the kinds of things that happen every day. They didn't transfer the money, obviously.

PROFESSOR: The email was totally believable?

PROFESSOR: Yeah.

PROFESSOR: I've seen it.

PROFESSOR: Even the tone and everything seemed very plausible.

PROFESSOR: Yep.

PROFESSOR: They actually used email messages this senior executive had written that were posted on public websites, because they send memos to the community and things like that that used the exact same style, introduction, the way they closed the

message. Even the language and terms they used was identical. This is stuff they had used in other ways. So it was actually kind of creepy, because even when I first read it, if you didn't know what it was, it was semi-plausible.

And thankfully, the staff member that it was requested of-- like this was kind of an out of the ordinary request, it got flagged, even though it looked legitimate. And he or she responded directly to the sender, removing the Reply To address and actually putting in a known good one from his address book or her address book, and responded and asked and kind of unveiled the whole scheme. But it could have gone south really quickly.

OK, so I mentioned that the network-based vector in my experience isn't as prevalent as it used to be. And I'll kind of belie that here, or make a lie of it here. This year has been a year of major exploits. Every single major SSL implementation has been targeted. There was Schannel on the Microsoft side. There was the Apple implementation, open SSL. There was Poodle with SSL v3.

This is SSL, right? This is a security service. So when you put a service up facing the world, you're going to run SSL. So we had a lot of world-facing services out there that were vulnerable to some of these things.

Shellshock was one that affected the Bash shell where you could remotely exploit a system. So these are all kind of the gold standard of a network-based exploit. They were remotely exploitable, and they could get administrative privileges. So it's been a kind of a nasty year, in terms of that. So it's, in my opinion, a bit of an outlier.

But how do we deal with this? Because these services are public. They need to be public. We can't just fence them off, because they're vulnerable to something.

The first thing is automatic patching. In the old days, the latency between a zero day coming out and a patch coming out was fairly long. That's shortening, and shortening, and shortening, so we're down to literally hours.

So when these things surface, we're able to push out updates to at least the systems that we maintain. The systems we don't maintain, we're able to use our

communication folks, like the communications office and Monique on the security team, to craft messages to go out to the community to at least alert them to the fact that you really need the patch, because this is dangerous and it's out there.

On the more active front, we can detect these scans. So the StealthWatch tool, I mentioned way back, is a tool that pulls NetFlow data off of our network devices. And we can do some basic heuristics on that. And if we see an outside IP address talking to several hundred MIT systems, that's probably not good.

It could be good. And if it's good, we will totally white list it. And we've done that many, many times for research projects and just things that are legitimate. But it's probably not a bad posture to say, OK, if we see that, it's probably bad. Let's block it.

So we actually have some automated BGP null-routing going on where we're actually watching the flows. And if we see an anomalous behavior, we null-route on-the-fly. That runs every five minutes. So as soon as a scan starts, we cut it off at the knees.

On a more proactive front even than that is we will proactively scan. So in the case of Shellshock and some of the earlier SSL vulnerabilities that were really deadly, we actually scanned the community and sent out lists to those we had contact info for to let them know, hey, this IP address is running a service. This is a known vulnerable. Please patch it. It's really just about getting the information out to the community as quickly as we can.

OK future trends, because we're running short on time, consumerization of IT. I call them future trends, but the future is now here at MIT and pretty much at any .edu. These are old things that we've been dealing with.

Bring your own device-- I mean, I've owned my own phone here and at other institutions in the edu space forever. It makes policy enforcement really difficult, because how do you enforce policy on a device that you didn't pay for that you don't manage?

Consumerization of services, here at MIT, we have an enterprise agreement with Dropbox now, so that you can store data up on Dropbox. Unlimited storage. That's open to students, right? So yes, unlimited storage on Dropbox, which is great.

The problems that come along with that are maybe data custody. Where is that data going to live? In our case, we've made sure that jurisdiction will always be in the United States. But what happens if you're dealing with a provider that crosses national boundaries into areas where they have different regulations?

What do you do if a person puts sensitive information up on that Dropbox and it gets synced up to the cloud. And they think it just lives in the cloud, but we know that Dropbox syncs to the local system. There's a lot of issues involved with the consumerization of IT, because the IT department doesn't control the service anymore. They're really just brokering the service between the service provider and the consumer.

Third-party email providers, kind of same thing. You might send sensitive information through an email system that's not totally internal. So sensitive data might leave the institution. Cloud-based resources kind of ties into that as well.

MIT never really had a perimeter. Neither does the rest of the world now, right? If you're a small startup-- and I'm sure you all have many friends at startups, I do as well-- none of them have local resources anymore, right? They're using stuff that's entirely in the cloud.

How do you draw a line or put a moat around those resources when they're living in Amazon Web Services, and at Salesforce.com, and as Google Apps, and as Dropbox? We need to find different ways to handle that.

We have the same data custody issues as to where that data might live. We've also got authentication and authorization issues. How do you make sure that just your users are accessing those services?

And that's where things like SAML come in, which I think MIT is really well-

positioned, because we have this really robust SAML architecture. When we wanted to add Dropbox, it was easy enough to add them as a service provider to our Touchstone infrastructure. And it just worked. I'm sure I'm glazing over some things, but you know, in the grand scheme, standards-based and federated systems, like SAML and Shibboleth are really life savers in a cloud-based world.

The internet of things, what does that mean? It seems to be the new buzzword du jour. But in terms of our experience so far, the internet of things at mit.net, we have building management systems all over, right? These are computer systems that are built by the fine folks that built air conditioners last year. So they're not all that secure, for the most part.

Mark had the story about they were just living on public mit.net. They could be probed by anyone in the world. What we've started to do-- and actually, we're almost entirely done with our building management systems-- is move them onto a different VRF, which is a Virtual Routing and Forwarding instance, so that they have a completely different forwarding path. And we front-end it with a firewall. It's all access controlled. It lives on separate physical infrastructure in closets. The closets are secured.

But when we move into a internet of things world, this problem is just going to multiply. What happens when the light switches have IP addresses and, who knows, my shoes have IP addresses? And it's going to get crazy. So how do we deal with that? And frankly, I don't have an answer quite yet.

Many companies say they do. And they'll make you spend a lot of money on solutions. But one of them I can think of is maybe we map access policy down to devices, based on 802.1X. So when I authenticate or my device authenticates, it pulls down the thermostat policy, so that it can coexist on the same network, and yet not be wide open to the world as, say, my laptop is.

So with that, I realized we're over, but are there any questions? Yes.

AUDIENCE: So there was one page you skipped over a coupld of slides ago.

PROFESSOR: Oh, sure.

AUDIENCE: Campus firewall.

PROFESSOR: Oh, I'm sorry. Where was this?

AUDIENCE: I'm curious about it.

PROFESSOR: [INAUDIBLE].

PROFESSOR: I swear, it wasn't intentional. Oh, yeah. Coming soon. So this is-- do you want to talk about it, or shall I?

PROFESSOR: Yeah. So I'll talk about it. I mean, one of the things we realized, as David said, is that your default posture for things, you get it with an Xbox you install it today. You have IP tables by default. Install Windows machine, you have Windows host firewall.

One of the things we look at is-- you know, you have this internet of things and this vast variety of devices on mit.net is having a more secure posture by default, so that devices, by default, may not necessarily be exposed to the entire public network.

And there are legitimate reasons people want to have a device on the public network. And that's fine. You know, one of the things that's great about MIT is, if people want to do that, you allow them to do that. They can do that in an automated way. They can do that by themselves. They don't need to ask a policy person. They don't need to do anything like that.

So what we're trying to move towards is really a network topology where, by default, people will be behind some layer of protection. If they want to go to a web page and enroll themselves in the public internet level of access, they can do that without talking to anyone. And it's automated. And it just happens within a minute or two.

And so I think what we're trying to do is just move the default security posture to something a little bit more secure, by default. But at the same time, we recognize that our goal here is to not really disrupt the innovative activities that happen here. And so if people want to do that, students or faculty, on an opt-in basis and go to

the web page, that's up to them. Any other questions?

AUDIENCE: I had one.

PROFESSOR: Sure.

AUDIENCE: What's the traffic like now on the MIT network? It's like, what kind of traffic do you see the most of?

PROFESSOR: Yeah. So I mean, looking at StealthWatch, I'd say like 80% of our traffic, if you look at it by protocol, is like HTTP, you know, [INAUDIBLE].

PROFESSOR: Which would include HTTP-streamed media.

PROFESSOR: Movies, media. Now, I think the interesting question you could ask is how much of it is legitimate research activity.

[LAUGHTER]

I know you guys are all studying hard. I know I was, that's why I'm still working here. No. But yeah, it's interesting breakdown. I think the one thing we do philosophically, as a provider, is-- you know, a lot of schools, there was a time where they were trying to make judgments about what kinds of traffic and how much were going across their campus networks. MIT does not do that. One thing we believe very strongly in is not me nor anybody else in the administration is in a position to pass value judgment over someone's internet activity.

PROFESSOR: Because people live here.

PROFESSOR: You live here, right?

PROFESSOR: It's not just your work. I mean, people are doing a lot of research on Netflix, because we have a Netflix cache. And we a lot of traffic going there. But we also have thousands of students and staff living here. So that's them at night powering up their Netflix box and streaming.

PROFESSOR: Or whatever it may be. So we've always been in the position of like, you know, we

have some very detailed information about what it is. But I'd say that most of it is, I'd say, even nowadays, is kind of scary. I'd say, at night, half of it's video stream, which is--

AUDIENCE: Would you allow torrent [INAUDIBLE]?

PROFESSOR: Porn?

PROFESSOR: We don't. No.

AUDIENCE: Torrents [INAUDIBLE].

PROFESSOR: Oh. You know what's interesting? So porn and torrents are pretty similar, so yeah.

[LAUGHTER]

Not that that's what went through my head Freudian-wise. I'm sorry. You know, it's interesting. Torrent traffic has actually gone down. I think that's what's been interesting. I'd say it's actually gone down over the years.

I think, on the plus side, most things are getting so easy for people to get through something like a Netflix, or an Amazon Prime, or whatever it is, where you can subscribe for $4 a month, where people are generally doing it. We have a Comcast video TV service we offer for free to the students. If you want to do IPTV, you can just do it on your computer now. But for the most part, I've seen torrents have actually gone down. I would just be honest, which is kind of a surprise.

PROFESSOR: All right. Well, let's thank Dave and Mark.

PROFESSOR: Thank you, guys.

PROFESSOR: Thank you.

[APPLAUSE]